

(12) UK Patent Application (19) GB (11) 2 273 588 (13) A

(43) Date of A Publication 22.06.1994

(21) Application No 9226598.2

(22) Date of Filing 21.12.1992

(71) Applicant(s)

Calder Limited

(Incorporated in Ireland)

Broad Street Centre, Broad Street, County Waterford,
Ireland

(72) Inventor(s)

Martin Bolger

(74) Agent and/or Address for Service

Marks & Clerk

57-60 Lincoln's Inn Fields, LONDON, WC2A 3LS,
United Kingdom

(51) INT CL⁵

G06F 12/00

(52) UK CL (Edition M)

G4A AMX

(56) Documents Cited

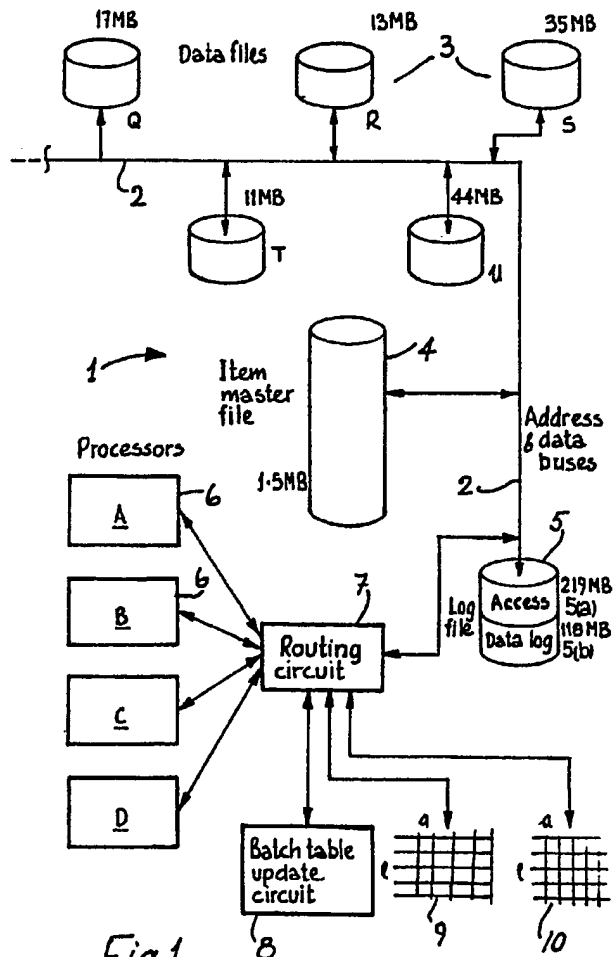
GB 2257273 A GB 2246221 A

(58) Field of Search

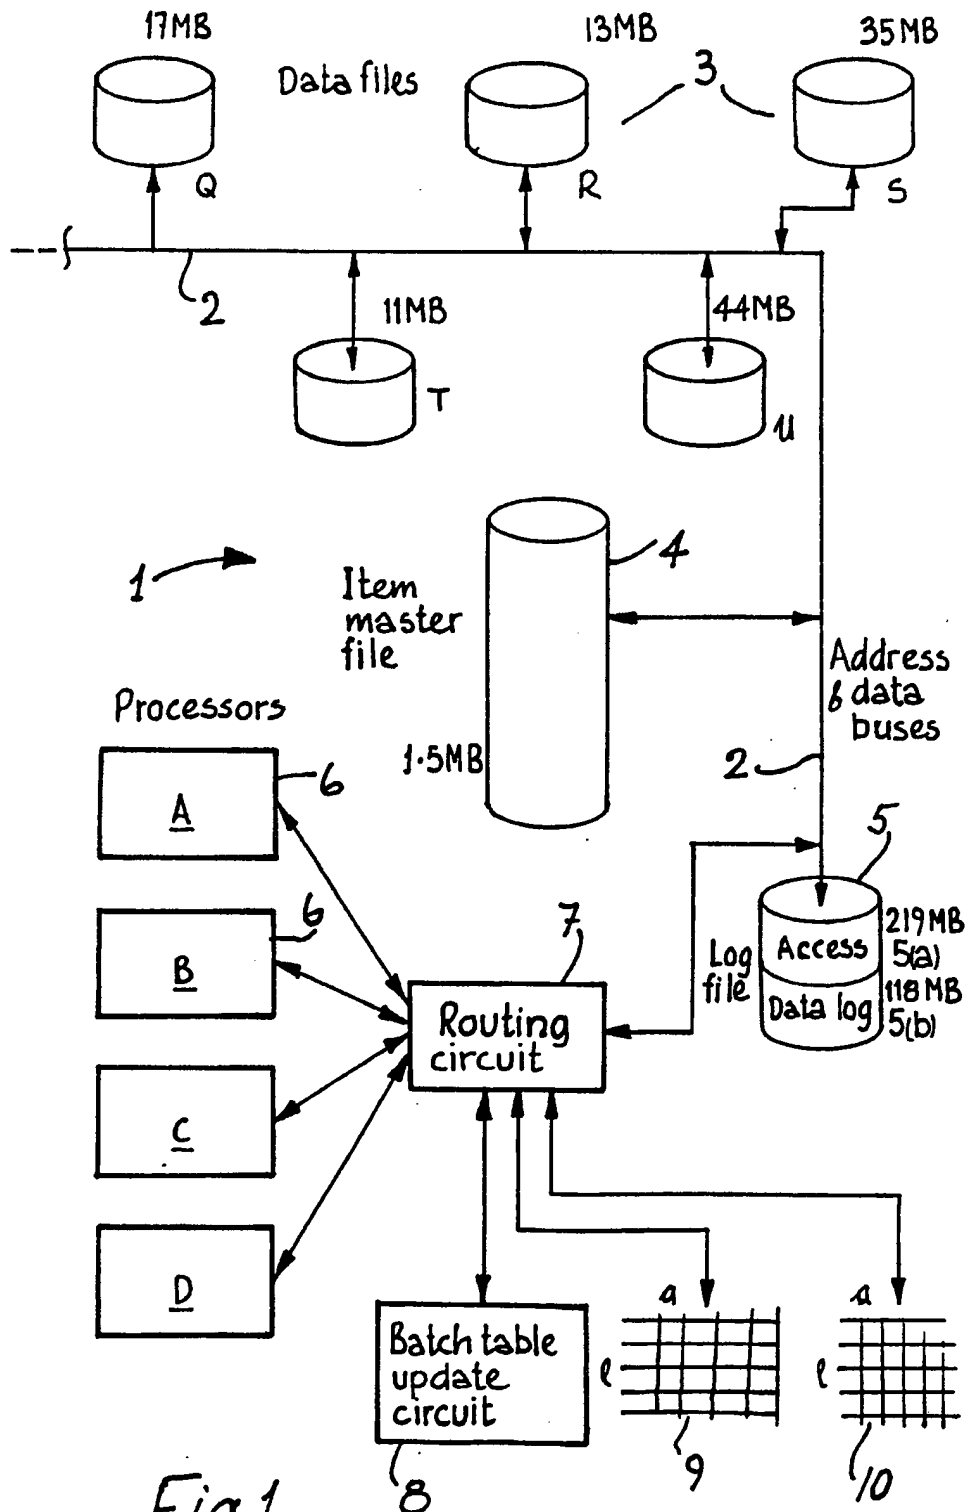
UK CL (Edition L) G4A AMG1 AMG2 AMX ANV
INT CL⁵ G06F 12/00 12/02 12/10

(54) Memory utilisation and data retrieval

(57) A memory mapping and data retrieval method involves setting up a number of data files 3 which have a header storing values for a number of parameters. An item master file 4 is also set up which has a record for each of the items to which the data relates. Logical addresses are used to determine actual addresses by reference to look-up tables 9, 10, one of which is updated on-line, the other being updated in a batch process. All updates to a data file 3 are recorded in a log file 5 which also includes a log of an access path to the stored data. Because the headers 3a have a number of parameter values, different processors 6 may access the same data file 3 using different logical addresses.



GB 2 273 588 A



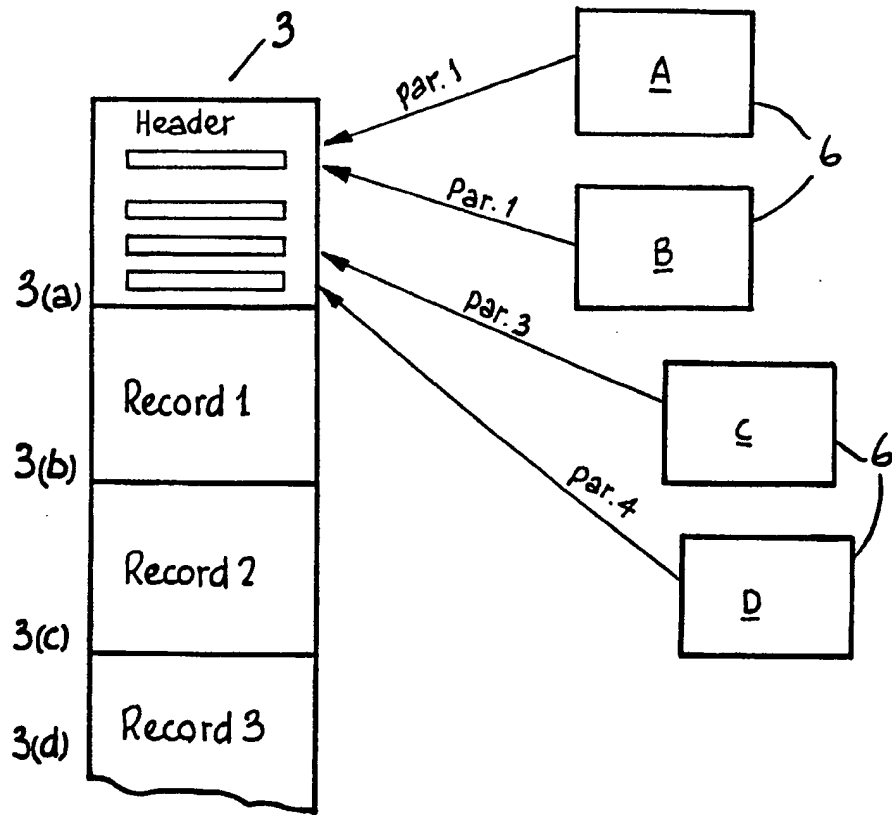


Fig. 2

"Memory Utilisation and Data Retrieval"

The invention relates to the utilisation of memory locations in a data processing system and to retrieval of data from these locations.

5 In more detail, the invention relates to the control of database operations for transaction data processing where there is a very large amount of data stored and fast access to the data is required.

British Patent Specification GB 2202656B discloses a method for memory re-mapping in a microcomputer system.
10 The present invention is directed towards providing a method of mapping memory locations and controlling access to the data in a large database system in which there may be in the order of 10GB of stored data.

15 One object of the invention is to provide for the storage of data in the memory in such a manner that it may be accessed quickly, and also to provide an access control method for the data. Another object is that data integrity be maintained at a high level. This is very important for transaction processing where there is a very large amount
20 of transactions occurring.

According to the invention, there is provided a memory utilisation and data retrieval method comprising the steps of:-

25 storing on a storage device an item master file having a record storing fixed data for each item to which the data relates;

storing on a storage device a set of data files, each storing data relating to a particular data process;

5

storing on a storage device a log file, comprising a data log for reception of transaction data and an access file for storing paths to data in the data log;

10

storing at least two look-up tables cross-referencing logical with actual address for data files and records of the item master file;

15

carrying out processing and retrieving data for the processing by accessing both the record of the item master file which relates to the relevant item and retrieving data from a data file which relates to a particular process, this step comprising the sub-steps of determining a logical address for the data files and for the record of the item master file by reference to one of the stored look-up tables; and

20

25

upon completion of an update operation to one of the data files, writing data relating to the update of the data log and subsequently storing in the access file of the log file the paths to the data in the data log.

30

In one embodiment, the logical address is one of a plurality of parameter values stored in a header of each data file so that different processors may have different logical addresses for the same data file.

Preferably, the method comprises the further step of automatically updating a look-up table upon completion of a data processing operation.

5 In another embodiment, another look-up table is updated in a batch process carried out at periodic time intervals.

The invention will be more clearly understood from the following description of some preferred embodiments thereof, given by way of example only but with reference to the accompanying drawings:-

10 Fig. 1 is a diagrammatic representation of a transaction data processing system of the invention; and

Fig. 2 is a detailed view showing the manner in which records are created.

15 Referring to the drawings, there is shown a transaction data processing system, indicated generally by the reference numeral 1. The system 1 comprises bi-directional address and data buses 2 which link up various processors and fixed disks storing database files. In
20 this embodiment there is a total fixed disk capacity of 8.3GB spread out between fourteen disks, which have an average utilisation in the order of 80-85%. Not all of the address and data buses are shown, for clarity. The fixed disks store approximately 800 data files (3) only
25 five of which are shown, namely, Q, R, S, T, U. The size of the data files are as follows:-

30 Q - 17MB
R - 13MB
S - 35MB
T - 11MB
U - 44MB.

The data files have a structure somewhat as shown in Fig. 2 wherein there is a header 3(a) and a number of records, three of which are shown, namely, 3(b), 3(c) and 3(d). The records are all of equal size. The header 3(a) is
5 described in more detail below.

Possibly the most important file stored on the fixed disks is an item master file 4 which is relatively small, being 1.5MB in size. The item master file 4 is generated at the beginning of database operations and it includes a data
10 record relating to each of a large number of items involved in the transactions. For example, in retail trading processing, each item would be a particular product. The data which is stored in each record is generally of fixed nature as it rarely changes. In this
15 embodiment, the data includes a status flag indicating whether or not a product is active, dormant (i.e. not being traded at present, but may be in the near future), or dead (i.e. no longer being traded). Other data includes an identifier for the product including a product number and
20 also a textual description of the product. While the items of this embodiment are products for retail trading, it is envisaged that the items may be of any other type, depending on the nature of the processing.

Another important file is a log file 5 which includes an
25 access file 5(a) and a data log 5(b). The access file has a size of 219MB, while the data log has a size of 118MB. The log file 5 is described in more detail below. The system 1 also includes a number of processors 6, namely, processors A, B, C, and D, all of which are connected to
30 a routing circuit 7. The routing circuit 7 is connected to a batch table update circuit and also to a pair of look-up tables 9 and 10.

In operation, the processors 6 operate to carry out various transaction operations using data stored in the data files 3 of the fixed disk. For retail trading processing, the processor A may carry out sales analysis, the processor B purchase orders, the processor C inventory calculations and the processor D generating management information reports. The nature of the processing is not important, the important point being that the processing involves accessing data from the fixed disks. Each time data is required, a request is transmitted to the routing circuit 7 which reads the logical address for the relevant data file in the instructions received from the relevant processor 6. The routing circuit 7 then accesses one or other of the look-up tables 9 and 10 to determine an actual address corresponding to the embedded logical address. The actual address is then transmitted on the address bus 2 and used to read the relevant item record from the item master file, and subsequently to retrieve the relevant data file 3. In more detail the logical address of the data file may be one of a number of keys which are stored in the header 3(a) of the data file 3(b). The header 3(a) stores values for four parameters. Each of the processors will have stored one of these parameter values as a logical address or primary key for that data file. As shown in Fig. 2 processors A and B both have the value of parameter 1 as the logical address, processor C has the value for parameter 3, whereas the processor D has the value for the parameter 4. Thus, all of the four processors may access the particular data file 3 shown in Fig. 2, however, different logical addresses are used in each case. This is extremely important in allowing versatility in the processing operations which are carried out and to avoid the need for duplication in storage of data files. The actual addresses for the files may be easily determined by reference to the look-up tables 9 and 10 and thus there is a minimum delay in operation of the

routing circuit 7. Further, because the relevant item record is always read from the item master file 4 a large amount of data may be read very quickly with a minimum amount of storage space being required.

5 The look-up tables 9 and 10 may be updated in two ways. The look-up table 9 is updated in-line by the routing circuit 7 according to instructions received from the processors 6. However, the look-up table 10 is only
10 updated at periodic time intervals under instructions from the batch table update circuit 8. In general, the look-up table 10 is only referred to by processors which operate in batch mode, after operation of the update circuit 8. It has been found that by segregating the
15 look-up tables into two different types, one being updated in-line and the other being updated in a batch process considerably reduces the amount of processing capacity required during transaction processing.

The processor 6 operates to write data to a new record in the log file 5 every time one of the data records 3 is
20 updated. Thus, at the end of a time period, there will be a record in the data log 5(b) of the log file relating to each update operation to a data file 3. Because of this, certain processes such as those for generation of documents may be carried out much more efficiently simply
25 by accessing the data log 5(b) using the access file 5(a) within the log file 5. This saves a considerable amount of fetch cycles to the data files 3 for certain types of processes. It also provides a method of monitoring database operations in an extremely efficient manner by
30 simply dumping data from the log file 5 to a report file whenever required. When retrieving data from the data log 5a the look-up tables 9 and 10 are not used, as the access file 5a provides the routes or paths and these are updated

on an on-going bases as data is being written to the data log 5(b).

5 It will be appreciated that the invention provides for the efficient utilisation of storage locations and to the provision of fast response times to data request.

The invention is not limited to the embodiments hereinbefore described but may be varied in construction and detail.

CLAIMS.

1. A memory utilisation and data retrieval method comprising the steps of:-

- 5 storing on a storage device an item master file having a record storing fixed data for each item to which the data relates;
- storing on a storage device a set of data files, each storing data relating to a particular data process;
- 10 storing on a storage device a log file, comprising a data log for reception of transaction data and an access file for storing paths to data in the data log;
- 15 storing at least two look-up tables cross-referencing logical with actual address for data files and records of the item master file;
- 20 carrying out processing and retrieving data for the processing by accessing both the record of the item master file which relates to the relevant item and retrieving data from a data file which relates to a particular process, this step comprising the sub-steps of determining a logical address for the data
- 25 files and for the record of the item master file by reference to one of the stored look-up tables; and
- upon completion of an update operation to one of the data files, writing data relating to

the update of the data log and subsequently storing in the access file of the log file the paths to the data in the data log.

2. A method as claimed in claim 1, wherein the
5 logical address is one of a plurality of parameter values stored in a header of each data file so that different processors may have different logical addresses for the same data file.
3. A method as claimed in claim 1 or 2 comprising the
10 further step of automatically updating a look-up table upon completion of a data processing operation.
4. A method as claimed in claim 3 wherein another
15 look-up table is updated in a batch process carried out at periodic time intervals.
5. A method substantially as hereinbefore described with reference to and as illustrated in the accompanying drawings.

Patents Act 1977
Examiner's report to the Comptroller under
Section 17 (The Search Report)

Application number

GB 9226598.2

Relevant Technical fields

(i) UK Cl (Edition L) G4A (AMG1, AMG2, AMX, ANV)

(ii) Int Cl (Edition 5) G06F (12/00, 12/02, 12/10)

Search Examiner

B G WESTERN

Databases (see over)

(i) UK Patent Office

(ii)

Date of Search

5 FEBRUARY 1993

Documents considered relevant following a search in respect of claims 1-5

Category (see over)	Identity of document and relevant passages	Relevant to claim(s)
E, A	GB 2257273 A (DEC) See whole document	1-5
A	GB 2246221 A (HITACHI) See whole document	1-5

Category	Identity of document and relevant passages	Relevance to claim

Categories of documents

X: Document indicating lack of novelty or of inventive step.

Y: Document indicating lack of inventive step if combined with one or more other documents of the same category.

A: Document indicating technological background and/or state of the art.

P: Document published on or after the declared priority date but before the filing date of the present application.

E: Patent document published on or after, but with priority date earlier than, the filing date of the present application.

&: Member of the same patent family, corresponding document.

Databases: The UK Patent Office database comprises classified collections of GB, EP, WO and US patent specifications as outlined periodically in the Official Journal (Patents). The on-line databases considered for search are also listed periodically in the Official Journal (Patents).